UT Austin Computer Sciences Technical Report AI07-348,
September 2007.

The UT Austin Villa 3D Simulation Soccer Team 2007

Shivaram Kalyanakrishnan and Peter Stone

Department of Computer Sciences
The University of Texas at Austin
Austin, TX 78712-0233
{shivaram, pstone}@cs.utexas.edu

Abstract. In this paper, we describe the soccer playing agent developed by our team UT Austin Villa to participate in the RoboCup 3D simulation soccer competition held in Atlanta, U.S.A., in July 2007. The 2007 competition saw the introduction of a humanoid robot in the 3D simulation league. We present an overview of this new version of the simulation league, followed by a description of our approaches in programming the humanoid robot. Our main emphasis was on developing skills such as walking and turning for the humanoid robot.

1 Introduction

In this paper, we describe the agent our team UT Austin Villa developed for participating in the RoboCup 3D Simulation Soccer competition held in Atlanta, U.S.A., in July 2007. The 3D simulation league has been in existence since 2004, but it underwent a significant change in 2007. While the simulated environment has remained a 3-dimensional world that models realistic physical forces such as friction and gravity, the soccer playing agent used until the 2006 competition was a sphere, which could be programmed through abstract commands such as **move** and **kick**. In keeping with the long-term goal espoused by the RoboCup federation, of pitting a team of 11 humanoid robots against 11 humans, it was decided to replace the sphere with a humanoid robot. The humanoid robot introduced in the 2007 competition requires programming at the level of control signals to its individual motors. With 20 degrees of freedom, it exposes the contestants to a problem whose nature and scale are dramatically different from one of programming the sphere through abstract, high-level actions.

Since the 2007 competition saw such a sea change from the previous years, and due to the limited time available to accommodate this change before the competition, the organizers and participants both were forced to quickly fix a set of specifications and work with them. But despite the pressing concerns of time and technical issues with the simulator, the 2007 competition turned out quite successful. Several teams were indeed able to exhibit agents with interesting and exciting behaviors. Also, the necessity of having a reasonably stable version of the simulator and a working set of game rules by the time of the competition ensured that several outstanding issues in the development were addressed, thereby streamlining the path for the future.

This paper is organized as follows. Section 2 provides a brief overview of the 3D humanoid simulator. In Section 3, we describe the design of the UT Austin Villa agent, and elaborate on its skills in Section 4. In Section 5, we draw conclusions and present directions for future work.

2 Brief Overview of 3D Simulation Soccer

In the 2007 competition, the humanoid robot in the simulation was derived from the Fujitsu HOAP-2 robot model [1]. The robot has 20 degrees of freedom: 6 in each leg and 4 in each hand. Figure 1 shows a visualization of the robot and the soccer field.

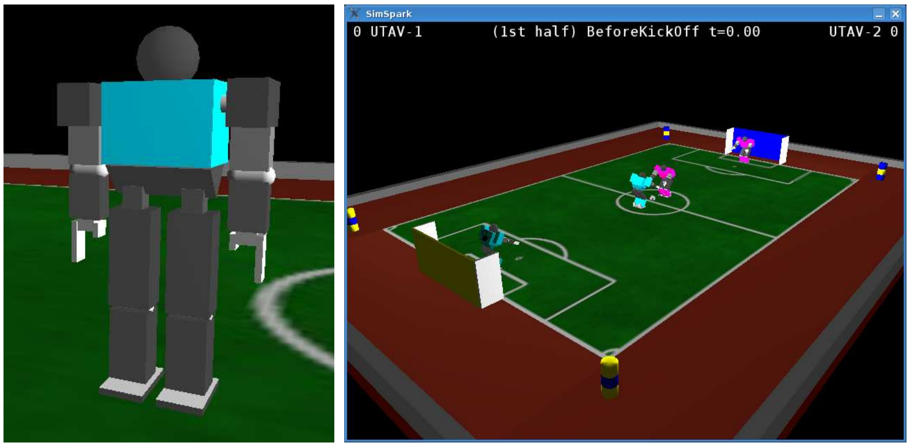

Figure 1. On the left is a screen shot of the robot, and on the right a view of a 2 versus 2 soccer game.

Each component of the robot's body is modeled as a rigid body with a mass that is connected to other components through joints. Torques may be applied to the motors controlling the joints. A physics simulator (Open Dynamics Engine [2]) computes the transition dynamics of the system taking into consideration the applied torques, forces of friction and gravity, collisions, etc. Sensation is available to the robot through a camera mounted in its torso, which provides information about the positions of all the objects on the field every cycle. In the 2007 competition, *noise-free* visual information was provided, with an *unrestricted* field of vision. The visual information, however, does not provide a complete description of state, as details like joint orientations of other players and the spin on the ball are not conveyed. Apart from the visual sensor, the agent also gets information from touch sensors at the feet, collision sensors, and gyro rate sensors. The simulation progresses in discrete time intervals with period 0.02 seconds. At each simulation step, the agent receives sensory information and is expected to return a 20-dimensional vector specifying torque values for the joint motors.

Since 2007 was the year the humanoid was introduced to the 3D simulation league, the primary focus was on developing robotic skills such as walking, turning, and kicking. This has itself been a challenging task, and is work still in progress. As a consequence, the 2007 competition did not present an ideal opportunity for the development of sophisticated high-level team behaviors. Interestingly, the team sizes were restricted to 2, and the games were conducted on a smaller version of the field than the standard; this was partly because existing computational resources were simply unable to cope with the requirements of running large teams of agents. Each game lasted 8 minutes; tie-breaking challenges included walking to the ball as quickly as possible, and taking a penalty kick.

The transition from the abstract "sphere" agent used until 2006 to the humanoid highlights the resolve of the RoboCup soccer community to embody intelligence in a realistic physical robot. The humanoid robot chosen for the task is indeed challenging to program; we expect that at least in the initial stages, significant time and effort will be devoted in developing low-level control algorithms and skills. Nonetheless, since this is a simulation league, the overhead involved in doing so compares very favorably with the time and labor entailed by a similar endeavor on a real humanoid robot. Indeed, we believe that the convenience afforded by simulation will enable the league to soon progress to tackle increasingly complex problems in humanoid robot behavior, possibly even ones that are yet to be realized in work with real humanoid robots. There are numerous vistas that research in the 3D humanoid simulation league is yet to explore; these provide the inspiration and driving force behind UT Austin Villa's desire to participate in this league.

3 Agent Architecture

At intervals of 0.02 seconds, the agent receives sensory information from the environment. The visual sensor provides distances and angles to different objects on the field from the agent's camera, which is located in its torso. It is relatively straightforward to build a world model by converting this information about the objects into Cartesian coordinates, particularly because visual sensation is complete and noise-free. Once a world model is built, the agent's control module is invoked. Figure 3 provides a schematic view of the control architecture of our humanoid soccer agent.

At the lowest level, the humanoid is controlled by specifying torques to each of its joints. We implement this through PID controllers for each joint, which take as input the desired angle of the joint and compute the appropriate torque. Further, we use routines describing inverse kinematics for the hands and legs. Given a target position and pose for the foot or the palm, our inverse kinematics routine uses trigonometry to calculate the angles for the different joints along the hand or the leg to achieve the specified target, if at all possible. The PID control and inverse kinematics routines are used as primitives to describe the agent's skills.

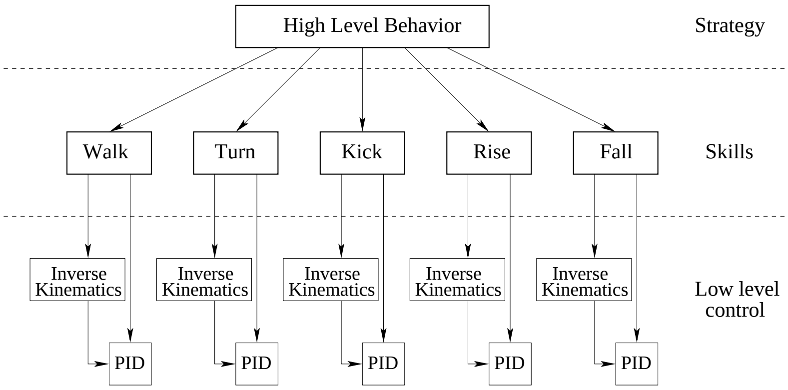

Figure 2. Schematic view of agent control architecture.

We found that the robot is much more stable while performing skills like walking and turning if it is bent slightly forward (as shown in Figure 3) than it is standing bolt upright. We maneuver the robot into such a position right at the start of the simulation, before it can start executing any of its skills. We elaborate on the agent's skills in section 4.

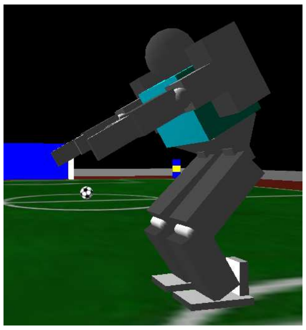

Figure 3. The humanoid is more stable when bent forward.

We did not attempt to develop very sophisticated high-level strategy for the team of agents, partly owing to the time constraints imposed by the competition, and partly because of the limitations presented by our set of skills. For our 2 versus 2 games, we instructed the goalie to remain standing a little in front of the goal, and to fall to either side to arrest a ball that is heading towards the goal. The other player always attempts to go to a position behind the ball and walk towards the opponent's goal, pushing (dribbling) the ball ahead of it. Ideally, we would have preferred to have the agent kick the ball towards the goal, but unfortunately, we were unable to implement a reliable kicking skill in time for the competition.

4 Player Skills

Our plan for developing the humanoid agent consisted of first developing a reliable set of skills, which could then be tied together by a module for high-level behavior. Needless to say, our foremost concern was locomotion. Bipedal locomotion is a well-studied problem (for example, by Pratt [8], and by Ramamoorthy and Kuipers [9]). However, it is hardly ever the case that approaches that work on one robot generalize in an easy and natural manner to others. Programming a bipedal walk for a robot demands careful consideration of the various constraints underlying it.

We experimented with several traditional approaches to program a walk for the humanoid robot, including monitoring its center of mass, specifying trajectories in space for its feet, etc. Through a process of trial and error, we concluded that a dynamically stable walk is more feasible than a statically stable walk (By "dynamically stable", we mean that the robot is balanced as long as it is moving; by "statically stable", we mean that the robot remains balanced even if it stops moving abruptly.). We achieved a reasonably fast dynamically stable walk by programming the robot to raise its left and right feet alternately (and perfectly out of phase) to a certain height above the ground and then stretching them back to their initial configurations. Inverse kinematics routines determine joint angles for the feet given the target position. As the robot beats its feet up and down, its crouched position causes it to translate forward. Interestingly, it is more appropriate to describe this walk routine as a "run", since there occur points in time during its execution when both feet lose contact with the ground.

By making minor changes to our walk routine, we were able to realize other useful skills for the robot. For getting the robot to turn, all that was required was to orient one of its hips at a slight angle while continuing to beat its feet up and down. Other slight variations of this basic pattern allowed for skills such as walking sideways and walking backwards. In fact, we were able to have the humanoid dribble the ball merely by walking "into" it; this was the only means to propel the ball forwards in the absence of a kick routine. To stop locomotion, the robot would gradually slow down the rate at which it took steps before coming to a complete halt.

Two other useful skills for the robot are falling and rising from a fallen position. We programmed the fall by having the robot bend its knee, by virtue of which it would lose balance and fall to one side. It is not surprising that we found it immensely more difficult to get a fallen robot to rise than to get a standing robot to fall. In our routine for rising, the fallen robot begins to thrash it arms and legs forward and backward until it rests on all fours with its head facing downwards. From this position, it slowly draws its feet closer to its arms, thereby bending significantly at the hips. When it stretches out again, the torso and hands are lifted up, and the robot returns to its upright position. The initial sequence of actions involving thrashing about is somewhat unnatural; however, it serves the purpose of maneuvering the robot into a downward-facing position. The ensuing sequence of actions to return the robot to its upright position strongly resembles a human-like routine for standing up.

Videos of our agent's skills are available at a supplementary web site [3].

5 Conclusions and Future Work

UT Austin Villa is pleased to be a part of the initiative of introducing a humanoid robot in the 3D simulation league. Despite our indifferent showing in the 2007 competition and the hurried nature of the competition itself, we believe that significant progress has been made all round from the point of view of research. The simulation of a humanoid robot opens up interesting problems for control, optimization, machine learning, and AI. The initial overhead for setting up the infrastructure is bound to be overtaken by the progress made through research on this important problem in the coming years. While the main emphasis thus far has been on getting a workable set of skills for the humanoid, it is conceivable that soon there will be a shift to higher level behaviors as well. A humanoid soccer league with scope for research at multiple layers in the architecture offers a unique challenge to the RoboCup community and augurs well for the future.

UT Austin Villa has been involved in the past in several research efforts involving RoboCup domains. Kohl and Stone [7] used policy gradient techniques to optimize the gait of an Aibo robot (4-legged league) for speed. Stone *et al.* [10] introduced Keepaway, a subtask in 2D simulation soccer [4, 6], as a test-bed for reinforcement learning, which has subsequently been researched extensively by others (for example, Taylor and Stone [11], Kalyanakrishnan *et al.* [5], and Taylor *et al.* [12]). We are keen to extend our research initiative to the 3D simulation league. Our initial focus will be on developing skills like kicking, which we are yet to deploy on our agent, and improve the skills that we have developed through the use of learning and optimization techniques. Banking on a reliable set of skills, we will seek to develop higher level behaviors like passing and intercepting.

References

1. Fujitsu Humanoid Robot HOAP-2. `http://jp.fujitsu.com/group/automation/en/services/humanoid-robot/hoap2/`.
2. Open Dynamics Engine. `http://www.ode.org/`.
3. UT Austin Villa 3D Simulation Team. `http://www.cs.utexas.edu/~AustinVilla/sim/3Dsimulation/`.
4. M. Chen, E. Foroughi, F. Heintz, Z. Huang, S. Kapetanakis, K. Kostiadis, J. Kummeneje, I. Noda, O. Obst, P. Riley, T. Steffens, Y. Wang, and X. Yin. Users manual: RoboCup soccer server — for soccer server version 7.07 and later. *The RoboCup Federation*, August 2002.
5. S. Kalyanakrishnan, Y. Liu, and P. Stone. Half field offense in RoboCup soccer: A multiagent reinforcement learning case study. *Proceedings of the RoboCup International Symposium 2006*, June 2006.
6. H. Kitano, M. Asada, Y. Kuniyoshi, I. Noda, E. Osawa, and H. Matsubara. RoboCup: A challenge problem for AI. *AI Magazine*, 18(1):73–85, 1997.
7. N. Kohl and P. Stone. Policy gradient reinforcement learning for fast quadrupedal locomotion. In *Proceedings of the IEEE International Conference on Robotics and Automation*, May 2004.
8. J. Pratt. *Exploiting Inherent Robustness and Natural Dynamics in the Control of Bipedal Walking Robots*. PhD thesis, Computer Science Department, Massachusetts Institute of Tehcnology, Cambridge, Massachusetts, 2000.
9. S. Ramamoorthy and B. Kuipers. Qualitative hybrid control of dynamic bipedal walking. *Robotics: Science and Systems*, II, 2007.
10. P. Stone, R. S. Sutton, and G. Kuhlmann. Reinforcement learning for RoboCup-soccer keepaway. *Adaptive Behavior*, 13(3):165–188, 2005.
11. M. E. Taylor and P. Stone. Behavior transfer for value-function-based reinforcement learning. pages 53–59, July 2005.
12. M. E. Taylor, S. Whiteson, and P. Stone. Comparing evolutionary and temporal difference methods for reinforcement learning. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 1321–28, July 2006.